Database System Providing Self–Tuned Parallel Database Recovery

Abstract

A database system providing self–tuned parallel database recovery is described. In one embodiment, for example, in a database system, a method is described for performing recovery operations using an optimal number of recovery threads, the method comprises steps of: (a) spawning an initial recovery thread to perform recovery operations; (b) measuring I/O (input/output) performance with the initial recovery thread; (c) spawning a subsequent recovery thread to perform recovery operations; (d) measuring I/O performance with the subsequent recovery thread; and (e) as long as I/O performance does not degrade beyond a preselected percentage, repeating steps (c) and (d) for spawning a desired number of additional recovery threads. In another embodiment, the database system auto–tunes the cache during performance of database recovery operations to optimize the performance of recovery operations.